arXiv:1704.02431v1 [cs.CV] 8 Apr 2017

Learning Cross-Modal Deep Representations for Robust Pedestrian Detection

Dan Xu[1], Wanli Ouyang[2,3], Elisa Ricci[4,5], Xiaogang Wang[2], Nicu Sebe[1]
[1]University of Trento, [2]The Chinese University of Hong Kong
[3]The University of Sydney, [4]Fondazione Bruno Kessler, [5]University of Perugia
{dan.xu, niculae.sebe}@unitn.it, eliricci@fbk.eu, {wlouyang, xgwang}@ee.cuhk.edu.hk

Abstract

This paper presents a novel method for detecting pedestrians under adverse illumination conditions. Our approach relies on a novel cross-modality learning framework and it is based on two main phases. First, given a multimodal dataset, a deep convolutional network is employed to learn a non-linear mapping, modeling the relations between RGB and thermal data. Then, the learned feature representations are transferred to a second deep network, which receives as input an RGB image and outputs the detection results. In this way, features which are both discriminative and robust to bad illumination conditions are learned. Importantly, at test time, only the second pipeline is considered and no thermal data are required. Our extensive evaluation demonstrates that the proposed approach outperforms the state-of-the-art on the challenging KAIST multispectral pedestrian dataset and it is competitive with previous methods on the popular Caltech dataset.

1. Introduction

Great strides in pedestrian detection research [3] have been made for challenging situations, such as cluttered background, substantial occlusions and tiny target appearance. As for many other computer vision tasks, in the last few years significant performance gains have been achieved thanks to approaches based on deep networks [21, 1, 17, 32]. Additionally, the adoption of novel sensors, *e.g.* thermal and depth cameras, has provided new opportunities, advancing the state-of-the-art on pedestrian detection by tackling problems such as adverse illumination conditions and occlusions [15, 11, 24]. However, the vast majority of wide camera networks in surveillance systems still employ traditional RGB sensors and detecting pedestrians in case of illumination variation, shadows, and low external light is still a challenging open issue.

This paper introduces a novel approach based on Convolutional Neural Networks (CNN) to address this problem. Our method is inspired by recent works demonstrating that learning deep representations from cross-modal data is greatly beneficial for detection and recognition tasks [12, 13]. However, most approaches assume the availability of large annotated datasets. In the specific case of pedestrian detection, the community can rely on a great abundance of visual data gathered with surveillance cameras, cars and robotic platforms, but there are few labeled multi-modal datasets. Therefore, motivated by the successes of recent unsupervised deep learning techniques, we introduce an approach for learning cross-modal representations for pedestrian detection which does not require pedestrian bounding box annotations. More specifically, we propose leveraging information from multispectral data and using a deep convolutional network to learn a non-linear mapping from RGB to thermal images without human supervision. This cross-modal mapping is then exploited by integrating the learned representations into a second deep architecture, operating on RGB data and effectively modeling multi-scale information. Importantly, at test time, thermal data are not needed and pedestrian detection is performed only on color images.

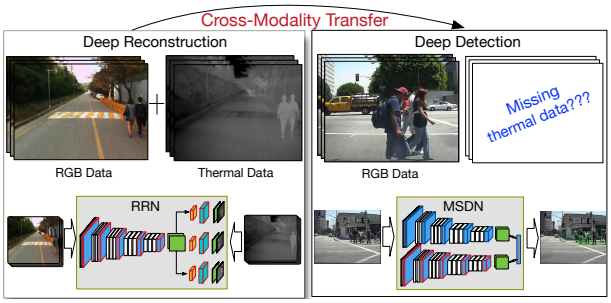

Figure 1. Overview of our framework. Our approach relies on two networks. The first network, named Region Reconstruction Network (RRN) is used to learn a non-linear feature mapping between RGB and thermal image pairs. Then, the learned model is transferred to a target domain where thermal inputs are no longer available and a second network, the Multi-Scale Detection Network (MDN), is used for learning an RGB-based pedestrian detector.

Figure 1 depicts an overview of the proposed approach. Our intuition, illustrated in Fig.2, is that, by exploiting multispectral data with the proposed method, it is easier to distinguish hard negative samples in color images (*e.g.*, electric

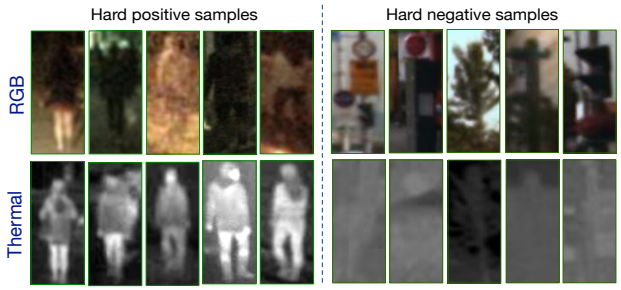

Figure 2. Motivation of this work. By exploiting thermal data in addition to RGB samples, it is easier to discriminate among pedestrians and background clutter.

poles or trees with appearance similar to pedestrians), thus improving the detection accuracy. Experimental results on publicly available datasets, where several frames are captured under bad illumination conditions, demonstrate the advantages of our approach over previous methods. To summarize the main contributions of this work are:

- We introduce a novel approach for learning and transferring cross-modal feature representations for pedestrian detection. With the proposed framework, data from the auxiliary modality (*i.e.* thermal data) are used as a form of supervision for learning CNN features from RGB images. There are two fundamental advantages in our strategy. First, multispectral data are not employed at the test phase. This is crucial when deploying robotics and surveillance systems, as only traditional cameras are needed, significantly decreasing costs. Second, no pedestrian annotations are required in the thermal domain. This greatly reduces human labeling efforts and permits to exploit large data collections of RGB-thermal image pairs.
- To our knowledge, this is the first work specifically addressing the problem of pedestrian detection under adverse illumination conditions with convolutional neural networks. Previous works mostly adopted hand-crafted descriptors and integrated the thermal modality by using additional input features [15, 28]. Our approach is based on two novel deep network architectures, specifically designed for unsupervised cross-modal feature learning and for effectively transferring the learned representations.
- Through an extensive experimental evaluation, we demonstrate that our framework outperforms the state-of-the-art on the novel KAIST multispectral pedestrian dataset [15] and it is competitive with previous methods on the popular Caltech dataset [9].

This paper is organized as follows. Section 2 outlines related work on pedestrian detection and cross-modal feature learning. Section 3 describes the proposed framework for learning features robust to illumination variations in the context of pedestrian detection. Experimental results to demonstrate the benefits of our approach are presented in Section 4. We conclude with key remarks in Section 5.

2. Related Work

Research topics closely related to this work are pedestrian detection from surveillance videos and deep learning approaches operating on multimodal data. Below, we present a review of the most recent works on these topics.

Pedestrian Detection. Due to its relevance in many fields, such as robotics and video surveillance, the problem of pedestrian detection has received considerable interests in the research community. Over the years, a large variety of features and algorithms have been proposed for improving detection systems, both with respect to speed [34, 2, 1, 17] and accuracy [39, 22, 44, 45, 10, 32].

Recently, notable performance gains have been achieved with the adoption of powerful deep networks [21, 1], thanks to their ability to learn discriminative features directly from raw pixels. In [26], a CNN pre-trained with an unsupervised method based on convolutional sparse coding was presented. The occlusion problem was addressed in [19], where a deep belief net was employed to learn the visibility masks for different body parts. This work was extended in [20] to model relations among multiple targets. More recently, in [31] DeepParts, a robust framework for handling severe occlusions, was presented. Differently from previous deep learning models addressing the occlusion problem, DeepParts does not rely on a single detector but it is based on multiple part detectors. Tian *et al.* [32] learned discriminative representations for pedestrian detection by considering semantic attributes of people and scenes. Cai *et al.* [4] introduced Complexity-Aware Cascade Training (CompACT), successfully integrating many heterogeneous features, both hand crafted and derived from CNNs. Zhang *et al.* [43] presented an approach based on the Region Proposal Network (RPN) [25] and boosted forests.

Other works focused on improving the computational times of CNN-based pedestrian detectors. For instance, Angelova *et al.* [1] proposed the DeepCascade method, *i.e.* a cascade of deep neural networks, and demonstrated a considerable gain in terms of detection speed. An in-depth analysis of different deep networks architectural choices for pedestrian detection was provided in [14]. To our knowledge, none of these previous works considers multi-modal data or tackles the problem of pedestrian detection under adverse illumination conditions.

Previous works have considered transferring information from other domains for constructing scene-specific pedestrian detectors. Wang *et al.* [35] proposed an unsupervised approach where target samples are collected by exploiting contextual cues, such as motions and scene geometry. Then, a pedestrian detector is built by re-weighting labeled source samples, *i.e.* by assigning more importance to samples more similar to target data. This approach was later extended in [42] to learn deep feature representations. Similarly, in [5] a sample selection scheme to reduce the discrepancy be-

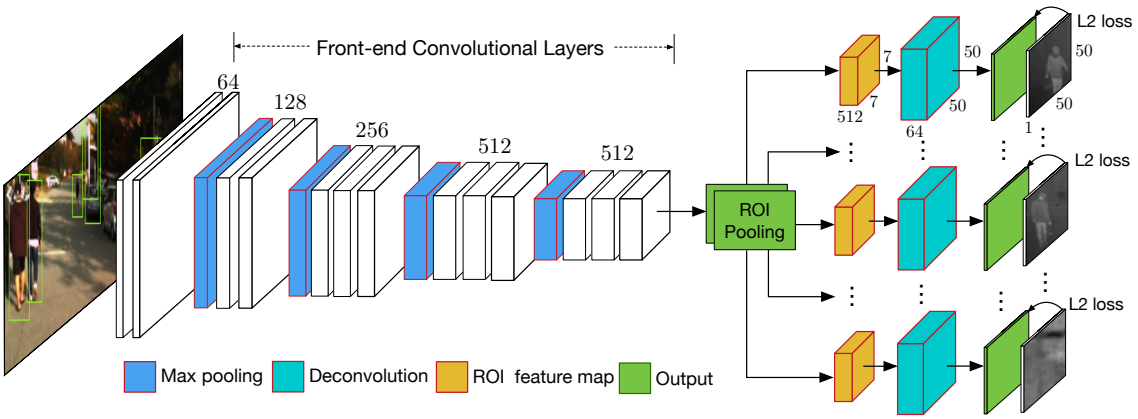

Figure 3. Architecture of the Region Reconstruction Network: a deep convolutional network trained for reconstructing thermal images from the associated RGB data. Best viewed in color.

tween source and target distributions was presented. Our approach is substantially different, as we do not restrict our attention to adapt a generic model to a specific scene and we tackle the problem of transferring knowledge among different modalities.

Learning Cross-modal Deep Representations. In the last few years deep networks have been successfully applied to learning feature representations from multi-modal data [16, 38, 37]. However, the problem of both learning and transferring cross-modal features has been rarely investigated. Notable exceptions are the works in [6, 30, 29, 12, 13]. Among these, the most similar to ours are [6, 30, 13]. In [6, 30] the idea of hallucinating data from other modalities was also exploited. However, our CNN-based approach is substantially different, since the work in [30] considered Deep Boltzmann Machines, while in [6] the mapping between different modalities was learned with Gaussian Processes. In [13] the problem of object detection from RGB data was addressed and depth images were used as additional information available only at training time. Similarly to [13], our detection network simultaneously use cross-modal features learned from a source domain and representations specific of the target scenario. However, in [13] labeled data were available in the original domain. Oppositely, in our framework we learn cross-modal features in an unsupervised setting, *i.e.* we do not require any annotation in the thermal domain. In this way, it is possible to exploit huge multispectral datasets.

3. Learning and transferring cross-modal deep representations

In this section we present the proposed framework. We first provide an overview of our approach and we describe in details the CNN architectures we design to reconstruct thermal data from RGB input and to transfer the learned cross-modal representations for the purpose of robust pedestrian detection.

3.1. Overview

As outlined in Section 1, the proposed framework (Fig.1) in based on two different convolutional neural networks, associated to the reconstruction and to the detection tasks, respectively. The first deep model, *i.e.* the Region Reconstruction Network (RRN), is a fully convolutional network trained on pedestrian proposals collected from RGB-thermal image pairs in an unsupervised manner. RRN is used to learn a non-linear mapping from the RGB channels to the thermal channel. In the target domain only RGB data are available and a second deep network, the Multi-Scale Detection Network (MSDN), embedding the parameters transferred from RRN, is used for robust pedestrian detection. MSDN takes a whole RGB image and a number of pedestrian proposals as input and outputs the detected bounding boxes with associated scores. In the test phase, detection is performed with MSDN and only RGB inputs are needed. In the following we describe the details of the proposed deepnet framework.

3.2. Region Reconstruction Network

The aim of RRN is to reconstruct thermal data from the associated RGB images. The design of the RRN architecture is driven by two main needs. First, in order to avoid human annotation efforts, thermal information should be recovered with an unsupervised approach. While our approach uses the thermal image as deep supervision for the reconstruction task, it essentially requires only very weak supervision information (*i.e.*, the pair-wise information). However, in the RGB-T data collection phase, we easily obtain the pair-wise information. The most expensive part in terms of human effort is to annotate the pedestrian bounding boxes. The proposed approach does not require these extra human-annotations. Second, as multispectral data are expected to be especially useful for hard positive and negative samples (Fig.2), instead of attempting to reconstruct the entire thermal images, it is more appropriate to specif-

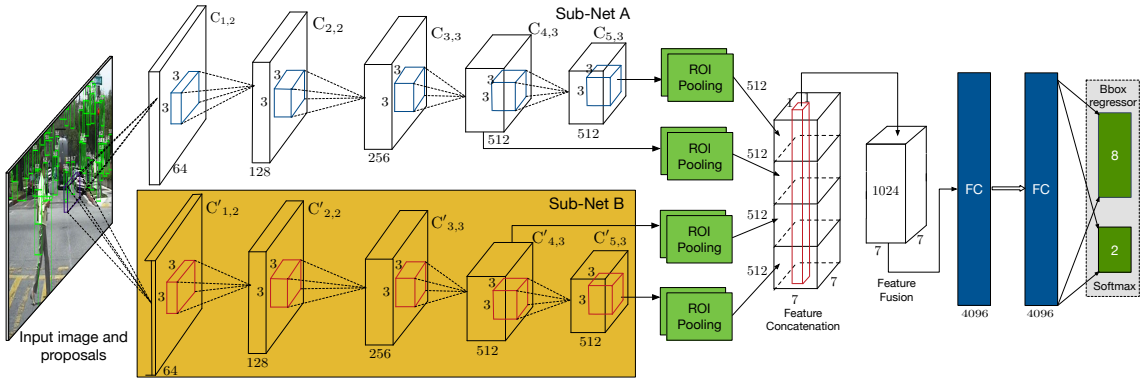

Figure 4. Architecture of the Multi-Scale Detection Network. Two sub-networks (Sub-Net A and Sub-Net B) with the same structure are used in MSDN. The parameters of all the convolutional layers of Sub-Net B (highlighted in yellow) are transferred from the Region Reconstruction Network.

ically focus on bounding boxes which are likely to contain pedestrians. Therefore, in this paper we propose to exploit a pretrained generic pedestrian detector (*e.g.* ACF [8]) to extract a set of pedestrian proposals (containing true positives and false positives) from RGB data and design a deep model which reconstructs the associated thermal information.

The proposed RRN network is illustrated in Fig.3. The input of RRN is a three-channel RGB image and a set of associated pedestrian proposals. RRN consists of a front-end convolutional subnetwork and a back-end reconstruction subnetwork. Although in our implementation the front-end convolutional layers exploit the VGG-13 network structure [27], RRN alternatively supports other architectures. After the last convolutional layer of the front-end subnetwork, an ROI pooling layer [10] is added. For each ROI, feature maps with size $512 \times 7 \times 7$ are generated. Considering the small size of the ROI feature maps, in order to effectively reconstruct the regions of thermal images associated to pedestrians, we apply a deconvolutional layer to upsample the ROI feature maps (output size 50×50) and reduce the number of output channels to 64 to ensure smooth convergence during training. Different from many previous works (*e.g.* [36]) which simply consider a bilinear upsampling operator, in the deconvolutional layer we learn the upsampling kernels (kernel size 4, stride 8 and pad 1). After the deconvolutional layer, a Rectified Linear Unit (ReLU) layer is applied. Then, reconstruction maps corresponding to each proposal are generated using a convolutional layer (kernel size 3, pad 1). Finally, a square loss is considered to compute each reconstruction map and the whole network is optimized with back-propagation.

In the widely used Fast- or Faster-RCNN frameworks, the groundtruth pedestrian bounding boxes are used to determine the ratio of true positive and false positive samples, and then construct fixed-size training mini-batches. To avoid using the carefully annotated groundtruth bounding boxes, we construct each training mini-batch using pedestrian proposals generated by thresholded generic ACF from one randomly selected training image, since the number of the proposals corresponding to each training image dynamically changes, our approach thus implements a dynamic mini-batch size during training.

3.3. Multi-Scale Detection Network

MSDN is specifically designed to perform pedestrian detection from RGB images by exploiting the cross-modal representations learned with RRN. Inspired by previous works demonstrating the importance of considering multi-scale information in pedestrian detection [43], we introduce a detection network which fuses multiple feature maps derived from ROI pooling layers.

MSDN architecture seamlessly integrates two sub-networks (Sub-Net A and Sub-Net B), as illustrated in Fig. 4. Sub-Net A has 13 convolutional layers, organized in five blocks. As depicted in Fig.4, $\mathrm{C}_{m,n}$ denotes the m-th block with n convolutional layers with the same size filters. Max pooling layers are added after the convolutional layers, and the ReLU non-linearity is applied to the output of each convolutional layer. An RoI (Region of Interest) pooling layer [10] is applied to the last two convolutional blocks to extract feature maps of size $512 \times 7 \times 7$ for each pedestrian proposal. We consider these two blocks, as our experiments show that this strategy represents the optimal trade-off between computational complexity and accuracy. Sub-Net B has the same structure of Sub-Net A but, since its main goal is to transfer cross-modality mid-level representations, the parameters of the 13 convolutional layers ($\mathrm{C'}_{1,2}$ to $\mathrm{C'}_{5,3}$) are derived from the associated layers of RRN. Indeed, the convolutional blocks from RRN produce a compact feature representation which captures the complex relationship among the RGB and the thermal domain. Therefore, they are embedded in MSDN, such as to allow the desired knowledge transfer.

The feature maps derived from the RoI pooling layers of the two sub-networks are then combined with a concatenation layer and a further convolutional layer with 1024

channels is applied. As the size of the RoI feature maps is small, we set the kernel size equal to 1 in the convolutional layer. Then, two fully connected layers of size 4096 follow. Finally, two sibling layers are used, one that outputs softmax probability estimates over pedestrian and background classes, and another that provides the associated bounding-box offset values for pedestrian localization.

3.4. Optimization

As discussed above, the proposed cross-modal framework is based on two different deep networks. Therefore, the training process also involves two main phases.

In the first phase, RRN is trained on multispectral data. The front-end convolutional layers of RRN are initialized using the parameters of the 13 convolutional layers of the VGG-16 model [27] pretrained on ImageNet dataset. The remaining parameters are randomly initialized. Stochastic Gradient Descent (SGD) is used to learn the network parameters. In the second phase, the parameters of MSDN are optimized using RGB data and pedestrian bounding box annotations in the target domain. We first train Sub-Net A by adding the common parts of MSDN (*i.e.* from the feature concatenation layer to the two sibling layers). In this case the size of the feature maps in the concatenation and in the following convolutional layers is $1024 \times 7 \times 7$ and $512 \times 7 \times 7$, respectively. The pretrained VGG-16 model is also utilized to initialize Sub-Net A. The convolutional layers of Sub-Net B are initialized with the corresponding parameters of RRN. Then, fine-tuning is performed using the RGB data of the target domain. The whole MSDN optimization is based on back-propogation with SGD.

3.5. Pedestrian detection

In the detection phase, given a test RGB image, we adopt the standard protocol. First, region proposals are extracted, similarly to the training phase. Then, the input image and the proposals are fed into MSDN. The softmax layer outputs the class score and the bounding box regressor indicates the estimated image coordinates. To reduce the redundancy of the proposals, non-maximum suppression is employed based on the prediction score of each proposal, setting an intersection over union (IoU) threshold δ.

4. Experiments

To evaluate the effectiveness of the proposed framework, we performed experiments on two publicly available datasets: the recent KAIST multispectral pedestrian dataset [15] and the popular Caltech pedestrian dataset [9]. In the following we describe the details of our evaluation.

4.1. Datasets

The **KAIST** multispectral pedestrian dataset [15] contains images captured under various traffic scenes with different illumination conditions (*i.e.* data recorded both during day and night). The dataset consists of 95,000 aligned RGB-thermal image pairs, of which 50,200 samples are used for training and the rest for testing. A total of 103,128 dense annotations corresponding to 1,182 unique pedestrians are available. We follow the protocol outlined in [15] in our experiments. The performance is evaluated on three different test sets, denoted as *Reasonable all*, *Reasonable day* and *Reasonable night*. *Reasonable* indicates that the pedestrians are not/partially occluded with more than 55 pixels height. The day and night sets are obtained from the *Reasonable all* set according to the capture time.

Figure 5. KAIST dataset. Reconstructed regions of thermal images (50×50 pixels) associated to the top nine detected pedestrian windows from ACF.

The **Caltech** pedestrian dataset [9] consists of about 10 hours of 30Hz video collected from a vehicle driving through urban traffic. The dataset contains 250,000 frames with 350,000 bounding boxes manually annotated and associated to about 2,300 unique pedestrians. Following previous works [32, 17], we strictly adopt the evaluation protocol in [9] measuring the log average miss rate over nine points ranging from 10^{-2} to 10^{0} False-Positive-Per-Image (FPPI). Our evaluation is conducted on both Caltech-All and Caltech-Reasonable settings.

Our approach uses RGB-thermal data for training, but in the test phase only requires RGB images as input. In all our experiments the KAIST training dataset is used to learn the RRN. Then, the performance of MSDN is assessed on the Caltech test set and on the RGB test frames of KAIST. The training and testing images of both datasets are resized (800 pixels height) to generate ROI feature maps with higher resolution useful for our reconstruction and detection tasks.

4.2. Experimental setup

Our framework is implemented under *Caffe*, and our evaluation is conducted on an Intel(R) Xeon(R) CPU E5-2630 with a single CPU core (2.40GHz), 64GB RAM and a NVIDIA Tesla K40 GPU.

We employ ACF [8] to generate pedestrian proposals for training both the reconstruction and the detection network with a low detection threshold of -70 as in [17] to obtain a high recall of pedestrian regions. In the test phase we also use ACF and consider the test proposals available online[1]. It is worth nothing that, while we focus on ACF, our cross-modality learning approach can be used in combination with an arbitrary proposal method.

For training the reconstruction network, we use the whole training set of the KAIST dataset. As thermal images captured from an infrared device have relatively low contrast and significant noise, we perform some basic processing, such as adaptive histogram equalization and denoising. By computing pedestrian proposals applying ACF, we end up creating a dataset of about 20K frames for training the region reconstruction network. All the frames are then horizontally flipped for data augmentation. We generate mini-batch of reconstruction RoIs from randomly chosen two images, and a fixed learning rate $\lambda_r = 10^{-9}$ is used to guarantee smooth convergence. We train the RRN for about 10 epochs.

For training the detection network on the the Caltech dataset we follow [45] and we construct a training set where every 3^{rd} frame is used. Instead, for the KAIST dataset we adopt the standard training protocol and every 20^{th} frame is considered. For both datasets, we use the same protocol for training MSDN. Similarly to RRN training, the data are flipped horizontally for the purpose of data augmentation. Each mini-batch consists of 128 pedestrian proposals randomly chosen from one training image. Positive samples with a ratio of 25% are taken from the proposals which have an IoU overlap with the ground truth of more than 0.5, while negative samples are obtained when the IoU overlap is in the range of [0, 0.5]. Stochastic gradient descent is used to optimize MSDN with the momentum and the weight decay parameters set to 0.9 and 0.0005, respectively. The network is trained for 8 epochs using an initial learning rate of 0.001 and drop by 10 times at the 5^{th} epoch.

4.3. Results on KAIST multispectral dataset

Analysis of proposed method. The first series of experiments aims to demonstrate the effectiveness of the proposed Cross-Modality Transfer CNN (CMT-CNN) framework. We evaluate the performance of our approach under four different settings: (i) CMT-CNN-SA. We only use Sub-Net A. The two ROI feature maps are concatenated and given as input to the convolutional fusion layer. This layer outputs a feature map with size 512, rather than 1024. Finally, the output is fed to the fully connected layers; (ii) CMT-CNN-SA-SB (ImageNet). We consider two sub-networks but initialize the convolutional layers of Sub-Net B from pretrained VGG16 model on ImageNet; (iii) CMT-CNN-SA-SB (Random): Same as (ii) but with random initialization for Sub-Net B; (v) CMT-CNN as described in Section 3, *i.e.* initializing the convolutional layers of Sub-Net B from trained RRN.

Table 1 shows the results of our comparison. Performance is evaluated using the log average miss-rate (MR). From the table it is clear that CMT-CNN significantly outperforms all its variations on all the three test sets, confirm-

[1] http://www.vision.caltech.edu/Image$_$Datasets/CaltechPedestrians/

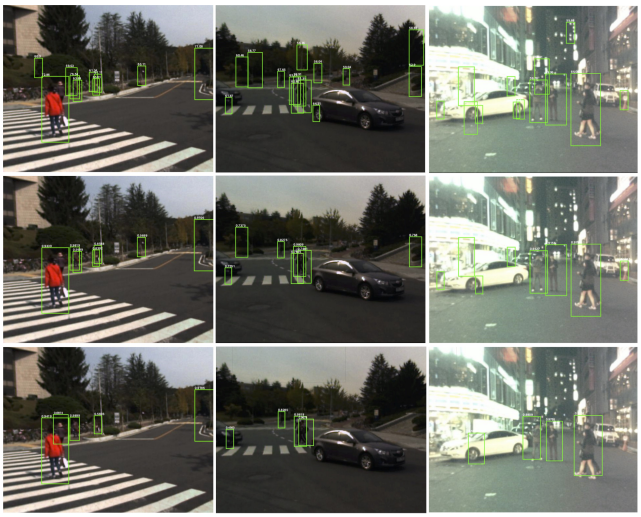

Figure 6. Examples of pedestrian detection results under different illumination conditions on the KAIST multispectral pedestrian dataset: (top) ACF detector, (middle) CMT-CNN-SA, (bottom) CMT-CNN.

Methods	All	Day	Night
CMT-CNN-SA	54.26%	52.44%	58.97%
CMT-CNN-SA-SB(Random)	56.76%	54.83%	61.24%
CMT-CNN-SA-SB(ImageNet)	52.15%	50.71%	57.65%
CMT-CNN	49.55%	47.30%	54.78%

Table 1. Comparison of different methods on the KAIST multispectral datasets including reasonable all, reasonable day and reasonable night settings.

Methods	Miss-Rate
CMT-CNN-SA	13.76%
CMT-CNN-SA-SB(Random)	15.89%
CMT-CNN-SA-SB(ImageNet)	13.01%
CMT-CNN-SA-SB(RGB-KAIST)	12.51%
CMT-CNN	10.69%

Table 2. Comparison of different variants of our method on the Caltech-Reasonable dataset. Performance are evaluated in terms of log-average miss-rate.

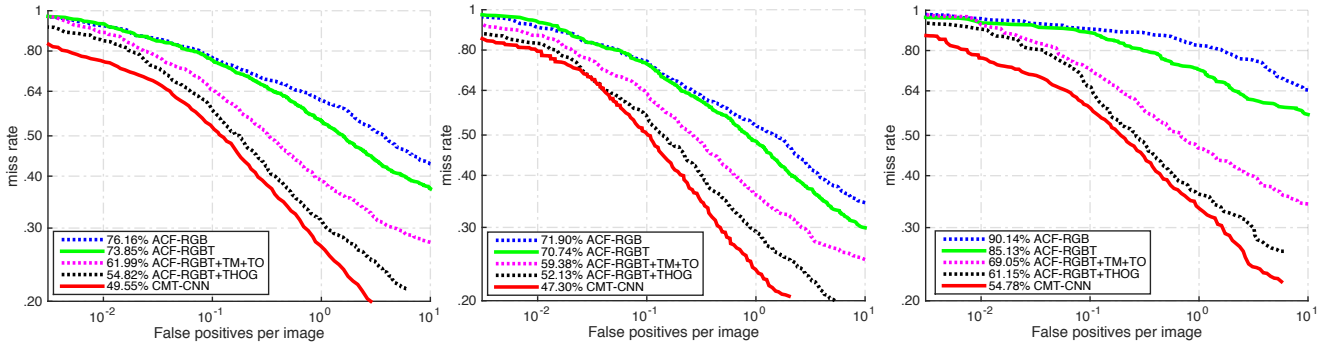

(a) Reasonable all (b) Reasonable day (c) Reasonable night

Figure 7. Quantitative evaluation results (miss rate versus false positive per image) on the KAIST multispectral dataset.

ing the fact that the proposed cross-modality framework improves the detection accuracy. We also observe that CMT-CNN provides lower MR than CMT-CNN-SA-SB, indicating that the performance gain of CMT-CNN is not only due to an increased number of parameters.

Figure 5 depicts some examples of the reconstruction results obtained with the proposed RRN. For the two given test frames, the reconstructed thermal regions associated to the top nine detection windows computed with ACF are shown. From the figure, it is easy to observe that the proposed network is able to effectively learn a mapping from RGB data to thermal data. Figure 6 shows some qualitative results obtained with MSDN. Comparing the detection bounding boxes of CMT-CNN-SA with those of CMT-CNN, we observe that hard negative samples are correctly classified with our method. For instance, the foliage from the trees (Fig. 6- first and second columns) is wrongly detected as pedestrian by CMT-CNN-SA. This confirms our intuition that leveraging information from multispectral data with our cross-modal representation transfer approach permits to improve the detection accuracy.

Comparison with state of the art methods. We also compare our approach with state of the art methods on the KAIST multispectral dataset. These methods include: (i) ACF-RGB [8], *i.e.* using ACF on RGB data; (ii) ACF-RGBT [15], *i.e.* using ACF on RGB-Thermal data; (iii) ACF-RGBT+TM+TO [15], *i.e.* using ACF on RGB-Thermal data with extra gradient magnitude and HOG of thermal images; (iv) ACF-RGBT+HOG [15], *i.e.* using ACF on RGB-Thermal data with HOG features with more gradient orientations than (iii). Results associated to these methods have been taken directly from the original paper [15]. Similarly to baseline approaches, we also use ACF for generating proposals both at training and at test time.

Observing Fig. 7, it is clear that CMT-CNN is several points better than ACF-RGBT+HOG, the best baseline on the KAIST dataset. Importantly, CMT-CNN only uses color images in the test phase, while ACF-RGBT+HOG exploits both RGB and thermal data. We also observe that on the *Reasonable night* setting, our approach obtains a more significant improvement than in the *Reasonable day* experiments. This demonstrates that CMT-CNN is especially useful for pedestrian detection under dark illumination conditions, thus confirming our initial intuition.

batch size	32	64	128	256
Caltech-All	65.97%	65.68%	65.32%	65.42%
Caltech-Reasonable	13.52%	13.01%	12.51%	12.35%

Table 3. Performance using different batch size in CMT-CNN-SA-SB (RGB-KAIST) experiments.

Method	Hardware	Miss-Rate	Testing Time (s/f)
InformedHaar [44]	CPU	75.85%	1.59
SpatialPooling [22]	CPU	74.04%	7.69
LDCF [18]	CPU	71.25%	0.60
CCF [40]	Titan Z GPU	66.73%	13.0
RPN + BF [43]	Tesla K40 GPU	64.66%	0.51
CompACT-Deep [4]	Tesla K40 GPU	64.44%	0.50
CMT-CNN	Tesla K40 GPU	64.01%	0.59

Table 4. Comparison of different methods (log-average miss-rate vs detection time). Log-average miss-rate is evaluated on the Caltech-All. s/f represents seconds per frame.

4.4. Results on Caltech pedestrian dataset

Analysis of CMT-CNN. Similarly to the experiments on the KAIST dataset, we first analyze the performance of our approach when different initialization strategies are used for Sub-Net B. In this case we also consider another baseline CMT-CNN-SA-SB (RGB-KAIST), *i.e.* we initialize Sub-Net B with VGG16 pretrained on ImageNet and further train it using RGB data of KAIST. The results of the comparison are shown in Table 2 and confirm the effectiveness of our framework. We observe that CMT-CNN-SA-SB (RGB-KAIST) beats CMT-CNN-SA-SB (ImageNet), showing that fine tuning CMT-CNN-SB with KAIST RGB data provides effective representations for improving the detection performance on Caltech. By using complementary data from the thermal modality, CMT-CNN further

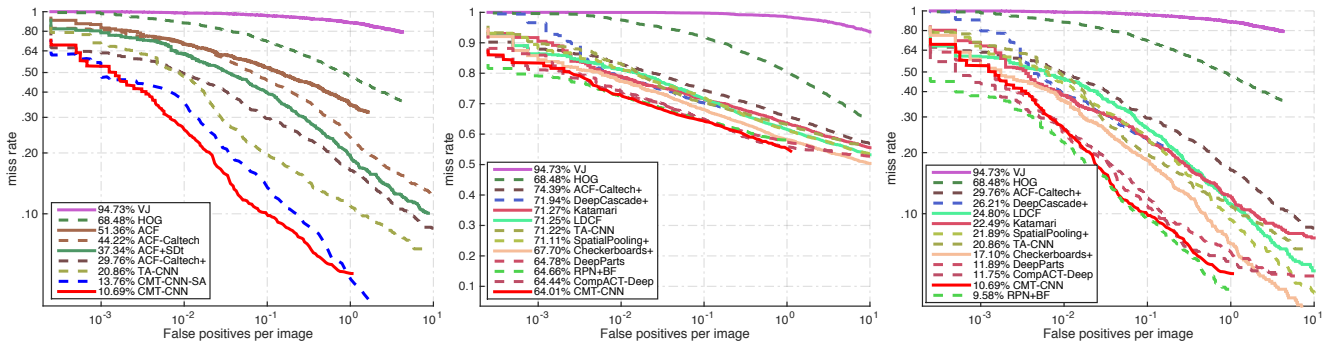

(a) (b) (c)

Figure 8. Quantitative evaluation results on the Caltech pedestrian dataset: comparison with (a) previous methods using ACF for proposals (VJ and HOG methods do not use ACF, but are kept as reference points) (b) state of the art methods on Caltech-All (c) state of the art methods on Caltech-Reasonable.

boosts its accuracy and outperforms CMT-CNN-SA-SB (RGB-KAIST). We observe that the improvement due to knowledge transfer on Caltech data is less pronounced than that obtained on KAIST dataset. We believe that this is mainly due to the fact that the frames of Caltech generally exhibit better illumination conditions than those of KAIST, while thermal information is especially beneficial in case of bad illumination.

To further demonstrate that the performance gain obtained with the proposed CMT-CNN is not simply due to ensembling different models, we consider the baseline CMT-CNN-SA-SB(RGB-KAIST) and we train Sub-Net B with KAIST RGB images using four different mini-batch size ranging from 32 to 256. For each experiment, the training samples are randomly shuffled. Table 3 shows the results of the four trials on Caltech-All and Caltech-Reasonable: using different batch size for Sub-Net B slightly affects the final performance and the best MR reported in the table is still worse than those obtained with CMT-CNN. This confirms the validity of our cross-modality learning approach.

We also compare the proposed CMT-CNN which uses ACF to generate region proposals with previous approaches also based on ACF proposals. Figure 8(a) shows the results of our comparison: our model outperforms all the baselines. Moreover, similarly to what we observed for KAIST experiments, CMT-CNN is more accurate than CMT-CNN-SA, confirming the advantage of our approach.

Comparison with state of the art methods. A comparison with state of the art methods is provided in Fig. 8(b). We considered Viola-Jones (VJ) [33], Histograms of Oriented Gradients (HOG) [7], DeepCascade+ [1], LDCF [18], SCF+AlexNet [14], Katamari [3], SpatialPooling+ [23], SCCPriors [41], TA-CNN [32], CCF and CCF+CF [40], Checkerboards and Checkerboards+ [45], DeepParts [31], CompACT-Deep [4] and RPN+BF[43]. Our approach attains a miss-rate of 10.69% on Caltech-Reasonable, which is very competitive with the state of the art methods, and a miss-rate of 64.01% on Caltech-All, which establishes a new state-of-the-art result. Importantly, our approach can be seen as complementary to most previous works. In fact, we believe that our unsupervised learning of cross-modal representations can be also integrated in other CNN architectures, to improve their robustness in coping with bad illumination conditions.

In Table 4 we report a comparison between our framework and recent pedestrian detection methods in terms of computational efficiency (times associated to previous methods are taken from the original papers). At test time, our network takes only 0.59 seconds to process one image, which is very competitive with previous methods.

5. Conclusions

We presented a novel approach for robust pedestrian detection under adverse illumination conditions. Inspired by previous works on multi-scale pedestrian detection [43], a novel deep model is introduced to learn discriminative feature representations from raw RGB images. Differently form previous methods, the proposed architecture integrates a sub-network, pre-trained on pairs of RGB and thermal images, such as to learn cross-modal feature representations. In this way, knowledge transfer from multispectral data is achieved and accurate detection is possible even in case of challenging illumination conditions. The effectiveness of the proposed approach is demonstrated with extensive experiments on publicly available benchmarks: the KAIST multispectral and the Caltech pedestrian detection datasets.

While this work specifically addresses the problem of pedestrian detection, the idea behind our cross-modality learning framework can be useful in other applications (*e.g.*, considering depth images for RGBD object/action detection and recognition). Hence, natural directions for future research include further investigating this possibility.

References

[1] A. Angelova, A. Krizhevsky, V. Vanhoucke, A. Ogale, and D. Ferguson. Real-time pedestrian detection with deep network cascades. In *BMVC*, 2015.

[2] R. Benenson, M. Mathias, R. Timofte, and L. Van Gool. Pedestrian detection at 100 frames per second. In *CVPR*, 2012.

[3] R. Benenson, M. Omran, J. Hosang, and B. Schiele. Ten years of pedestrian detection, what have we learned? In *ECCVW*, 2014.

[4] Z. Cai, M. Saberian, and N. Vasconcelos. Learning complexity-aware cascades for deep pedestrian detection. In *ICCV*, 2015.

[5] X. Cao, Z. Wang, P. Yan, and X. Li. Transfer learning for pedestrian detection. *Neurocomputing*, 100:51–57, 2013.

[6] C. M. Christoudias, R. Urtasun, M. Salzmann, and T. Darrell. Learning to recognize objects from unseen modalities. In *ECCV*, 2010.

[7] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005.

[8] P. Dollár, R. Appel, S. Belongie, and P. Perona. Fast feature pyramids for object detection. *TPAMI*, 36(8):1532–1545, 2014.

[9] P. Dollár, C. Wojek, B. Schiele, and P. Perona. Pedestrian detection: A benchmark. In *CVPR*, 2009.

[10] R. Girshick. Fast r-cnn. In *ICCV*, 2015.

[11] A. González, Z. Fang, Y. Socarras, J. Serrat, D. Vázquez, J. Xu, and A. M. López. Pedestrian detection at day/night time with visible and fir cameras: A comparison. *Sensors*, 16(6):820, 2016.

[12] S. Gupta, J. Hoffman, and J. Malik. Cross modal distillation for supervision transfer. In *CVPR*, 2016.

[13] J. Hoffman, S. Gupta, and T. Darrell. Learning with side information through modality hallucination. In *CVPR*, 2016.

[14] J. Hosang, M. Omran, R. Benenson, and B. Schiele. Taking a deeper look at pedestrians. In *CVPR*, 2015.

[15] S. Hwang, J. Park, N. Kim, Y. Choi, and I. So Kweon. Multispectral pedestrian detection: Benchmark dataset and baseline. In *CVPR*, 2015.

[16] A. Karpathy, A. Joulin, and F. F. Li. Deep fragment embeddings for bidirectional image sentence mapping. In *NIPS*, 2014.

[17] J. Li, X. Liang, S. Shen, T. Xu, and S. Yan. Scale-aware fast r-cnn for pedestrian detection. *arXiv preprint arXiv:1510.08160*, 2015.

[18] W. Nam, P. Dollár, and J. H. Han. Local decorrelation for improved pedestrian detection. In *NIPS*, 2014.

[19] W. Ouyang and X. Wang. A discriminative deep model for pedestrian detection with occlusion handling. In *CVPR*, 2012.

[20] W. Ouyang and X. Wang. Single-pedestrian detection aided by multi-pedestrian detection. In *CVPR*, 2013.

[21] W. Ouyang, X. Zeng, and X. Wang. Modeling mutual visibility relationship in pedestrian detection. In *CVPR*, 2013.

[22] S. Paisitkriangkrai, C. Shen, and A. van den Hengel. Strengthening the effectiveness of pedestrian detection with spatially pooled features. In *ECCV*, 2014.

[23] S. Paisitkriangkrai, C. Shen, and A. van den Hengel. Pedestrian detection with spatially pooled features and structured ensemble learning. *TPAMI*, PP(99):1–1, 2015.

[24] C. Premebida, J. Carreira, J. Batista, and U. Nunes. Pedestrian detection combining rgb and dense lidar data. In *IROS*, 2014.

[25] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.

[26] P. Sermanet, K. Kavukcuoglu, S. Chintala, and Y. LeCun. Pedestrian detection with unsupervised multi-stage feature learning. In *CVPR*, 2013.

[27] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.

[28] Y. Socarrás, S. Ramos, D. Vázquez, A. M. López, and T. Gevers. Adapting pedestrian detection from synthetic to far infrared images. In *ICCVW*, 2011.

[29] R. Socher, M. Ganjoo, C. D. Manning, and A. Ng. Zero-shot learning through cross-modal transfer. In *NIPS*, 2013.

[30] N. Srivastava and R. R. Salakhutdinov. Multimodal learning with deep boltzmann machines. In *NIPS*, 2012.

[31] Y. Tian, P. Luo, X. Wang, and X. Tang. Deep learning strong parts for pedestrian detection. In *ICCV*, 2015.

[32] Y. Tian, P. Luo, X. Wang, and X. Tang. Pedestrian detection aided by deep learning semantic tasks. In *CVPR*, 2015.

[33] P. Viola and M. J. Jones. Robust real-time face detection. *IJCV*, 57(2):137–154, 2004.

[34] P. Viola, M. J. Jones, and D. Snow. Detecting pedestrians using patterns of motion and appearance. *IJCV*, 63(2):153–161, 2005.

[35] M. Wang, W. Li, and X. Wang. Transferring a generic pedestrian detector towards specific scenes. In *CVPR*, 2012.

[36] S. Xie and Z. Tu. Holistically-nested edge detection. In *ICCV*, 2015.

[37] D. Xu, E. Ricci, Y. Yan, J. Song, and N. Sebe. Learning deep representations of appearance and motion for anomalous event detection. In *BMVC*, 2015.

[38] D. Xu, Y. Yan, E. Ricci, and N. Sebe. Detecting anomalous events in videos by learning deep representations of appearance and motion. *CVIU*, 156:117–127, 2017.

[39] J. Yan, X. Zhang, Z. Lei, S. Liao, and S. Li. Robust multi-resolution pedestrian detection in traffic scenes. In *CVPR*, 2013.

[40] B. Yang, J. Yan, Z. Lei, and S. Z. Li. Convolutional channel features. In *ICCV*, 2015.

[41] Y. Yang, Z. Wang, and F. Wu. Exploring prior knowledge for pedestrian detection. In *BMVC*, 2015.

[42] X. Zeng, W. Ouyang, M. Wang, and X. Wang. Deep learning of scene-specific classifier for pedestrian detection. In *ECCV*. 2014.

[43] L. Zhang, L. Lin, X. Liang, and K. He. Is faster r-cnn doing well for pedestrian detection? In *ECCV*, 2016.

[44] S. Zhang, C. Bauckhage, and A. Cremers. Informed haar-like features improve pedestrian detection. In *CVPR*, 2014.

[45] S. Zhang, R. Benenson, and B. Schiele. Filtered channel features for pedestrian detection. In *CVPR*, 2015.